



# NASA Advanced Computing Environment for Science & Engineering


Dr. Rupak Biswas

*Director, Exploration Technology Directorate  
Manager, High End Computing Capability (HECC) Project*


NASA Ames Research Center, Moffett Field, Calif., USA

15 March 2016

# NASA Overview: Mission Directorates






- **Vision:** *To reach for new heights and reveal the unknown so that what we do and learn will benefit all humankind*
- **Mission:** *To pioneer the future in space exploration, scientific discovery, and aeronautics research*
- **Aeronautics Research (ARMD):** Pioneer and prove new flight technologies for safer, more secure, efficient, and environmentally friendly air transportation
- **Human Exploration and Operations (HEOMD):** Focus on ISS operations; and develop new spacecraft and other capabilities for affordable, sustainable exploration beyond low Earth orbit
- **Science (SCMD):** Explore the Earth, solar system, and universe beyond; chart best route for discovery; and reap the benefits of Earth and space exploration for society
- **Space Technology (STMD):** Rapidly develop, demonstrate, and infuse revolutionary, high-payoff technologies through collaborative partnerships, expanding the boundaries of aerospace enterprise






# NASA Overview: Centers & Facilities




# Need for Advanced Computing



## Enables modeling, simulation, analysis, and decision-making

- Digital experiments and physical experiments are tradable
- Physical systems and live tests are generally expensive & dangerous (e.g., extreme environments), require long wait times, and offer limited sensor data
- NASA collects and curates vast amounts of observational science data that require extensive analysis and innovative analytics to advance our understanding




- Decades of exponentially advancing computing technology has enabled dramatic improvements in cost, speed, and accuracy – in addition to providing a predictive capability
- Many problems pose extremely difficult combinatorial optimization challenges that can only be solved accurately using advanced technologies such as quantum computing
- NASA's goals in aeronautics, Earth and space sciences, and human and robotic exploration all require orders-of-magnitude increase in computing capability to enhance accuracy, reduce cost, mitigate risk, accelerate R&D, and heighten societal impact




# Advanced Computing Environment




# NASA's Diverse HPC Requirements



- Engineering requires HPC resources that can process large ensembles of moderate-scale computations to efficiently explore design space (**high throughput / capacity**)
- Research requires HPC resources that can handle high-fidelity long-running large-scale computations to advance theoretical understanding (**leadership / capability**)




Aerodynamic database generation



Rotary wing unsteady aerodynamics

- Time-sensitive mission-critical applications require HPC resources on demand (**high availability / maintain readiness**)



Real-time hurricane prediction



# Balanced HPC Environment



## Computing Systems

- **Pleiades**: 212K-core SGI Altix ICE with 4 generations of Intel Xeon (4 racks GPU-enhanced: M2090, K40; 16 nodes have Phi 5110P); 723 TB RAM; 5.3 PF peak
- **Merope**: 12K-core SGI Altix ICE with 2 generations of Intel Xeon; 28 TB RAM; 141 TF peak
- **Endeavour**: Two SGI UV2000 nodes with 2 and 4 TB shared memory SSI via NUMALink-6; 32 TF peak
- **hyperwall**: 1024-core AMD Opteron, 128-node GPU M2090 cluster for large-scale rendering & concurrent visualization



## Data Storage


- 20 PB of RAID over several Lustre filesystems
- 115 PB of tape archive

## Networks

- InfiniBand interconnect for Pleiades in partial hypercube topology; connects all other HPC components as well
- 10 Gb/s external peering



# Integrated Spiral Support Services



## NASA Mission Challenges

Scientists and engineers plan computational analyses, selecting the best-suited codes to address NASA's complex mission challenges

**Outcome:** Dramatically enhanced understanding and insight, accelerated science and engineering, and increased mission safety and performance

## Performance Optimization



NAS software experts utilize tools to parallelize and optimize codes, dramatically increasing simulation performance while decreasing turn-around time


## Data Analysis and Visualization



NAS visualization experts apply advanced data analysis and rendering techniques to help users explore and understand large, complex computational results

## Computational Modeling, Simulation, & Analysis

NAS support staff help users to productively utilize NASA's supercomputing environment (hardware, software, networks, and storage) to rapidly solve large computational problems





# Accelerator Technologies



## Significant performance potential for science and engineering applications

- Execute many threads simultaneously at relatively lower power

## Two primary viable options


- Nvidia GPGPU: Did not get much traction within NASA
- Intel MIC: Code commonality across host and co-processor was initially promising

## Intel Xeon Phi (KNC) evaluation

- 128 nodes, each with 2 Sandy Bridge and 2 KNC
- Examine performance in four different execution modes: Native Host, Off-load, Symmetric, Native MIC
- *Micro-kernel benchmarks*: Memory bandwidth / latency, MPI functions, OpenMP constructs
- *NAS Parallel Benchmarks (NPB)*: OpenMP, MPI, MPI+OpenMP
- *Applications*: OVERFLOW, Cart3D, WRF
- Results reported without extensive code modifications



# Summary Performance Results



- System stability initially an issue but situation improved as MPSS (Many-core Platform Software Stack) has matured
- Running codes in Native modes lead to wasted resources
- MPI and OpenMP overhead very high on MIC compared to on host
- Off-load mode has significant overhead associated with data transfer
- Optimal load balancing in Symmetric mode is extremely challenging
- Hybrid code in Symmetric mode yields best performance due to reduced MPI communication and improved resource utilization
- Obtaining good performance on KNC is not simple – requires careful design of data structure and memory layout, and lots of parallelism
- KNC not ready for prime time, but next generation KNL looks promising due to no host and several other architectural improvements
- Extensive details in SC2013 paper by S. Saini et al.: *“An early performance evaluation of many integrated core architecture based SGI rackable computing system”*



# Big Data



## NASA has enormous collections of observational and model data

- Observational Data
  - Tens of satellites and telescopes producing multi-petabytes of data per year
  - SMD's Earth Science Division operates 12 DAACs (archive centers) containing ~10 PB of data
  - Solar Dynamics Observatory (SDO) satellite produces 1 GB per minute; translates to ~3 PB over its 5-year life cycle
- Model / Simulation Data
  - NAS Division has 20 PB of unique data in global filesystems and 115 PB of archive storage
  - MITgcm code running at 1/48<sup>th</sup> degree resolution on 35K cores produced 1.4 PB during its 5-day run; full simulation will produce 9–18 PB



## DISE (Data Intensive Supercomputing Environment) integrates Big Data and Big Compute to support analysis and analytics of NASA data

*Fun Fact:* The term “Big Data” was first used by Michael Cox & David Ellsworth of NAS Division in a Visualization ‘97 paper: “Visualizing flow around an airframe”, where largest dataset considered was 7.5 GB



# Big Data Challenges for Users

Conducted survey of NASA projects dealing with Big Data to gather user requirements

**BIG  
DATA**

Developing a roadmap including prototype implementations



- **Data Discovery:** Finding what data is available and where
- **Data Management:** Transferring very large datasets from archives to computational resources
- **Tools / Models / Algorithms:** Developing analysis & analytics software at scale
- **Analysis Workflow:** Handling increasingly complex processing pipelines
- **Analysis / Analytics Infrastructure:** Dealing with inadequacy of available heterogeneous resources
- **Collaboration Environments:** Difficulty with sharing knowledge across a wider community



# NASA Earth Exchange (NEX)




A collaborative environment that brings scientists and researchers together in a knowledge-based social network along with observational data, tools, and computing power to provide transparency and accelerate research



## VISION


To provide  
"Science as a  
Service" to the  
Earth science  
community  
addressing  
global  
environmental  
challenges



## GOAL

To improve  
efficiency and  
expand the  
scope of NASA  
Earth science  
technology,  
research, and  
applications  
programs

# NEX Environment



## Collaboration Portal (over 400 members)

Tools, Models,  
Workflows, Papers, Data



## Compute Resources (over 5 PF computing)

Sandbox for  
experimentation,  
HPC (Pleiades)



## Data Repository (over 2 PB of data)

Working copies of  
Observational and  
Project data

- Collaboration portal open to all Earth scientists
- Sandbox currently available to a subset of scientists with NASA credentials
- HPC resources available only to approved projects with allocation
- OpenNEX, a collaboration with Amazon, provides NEX datasets to the wider Earth science community

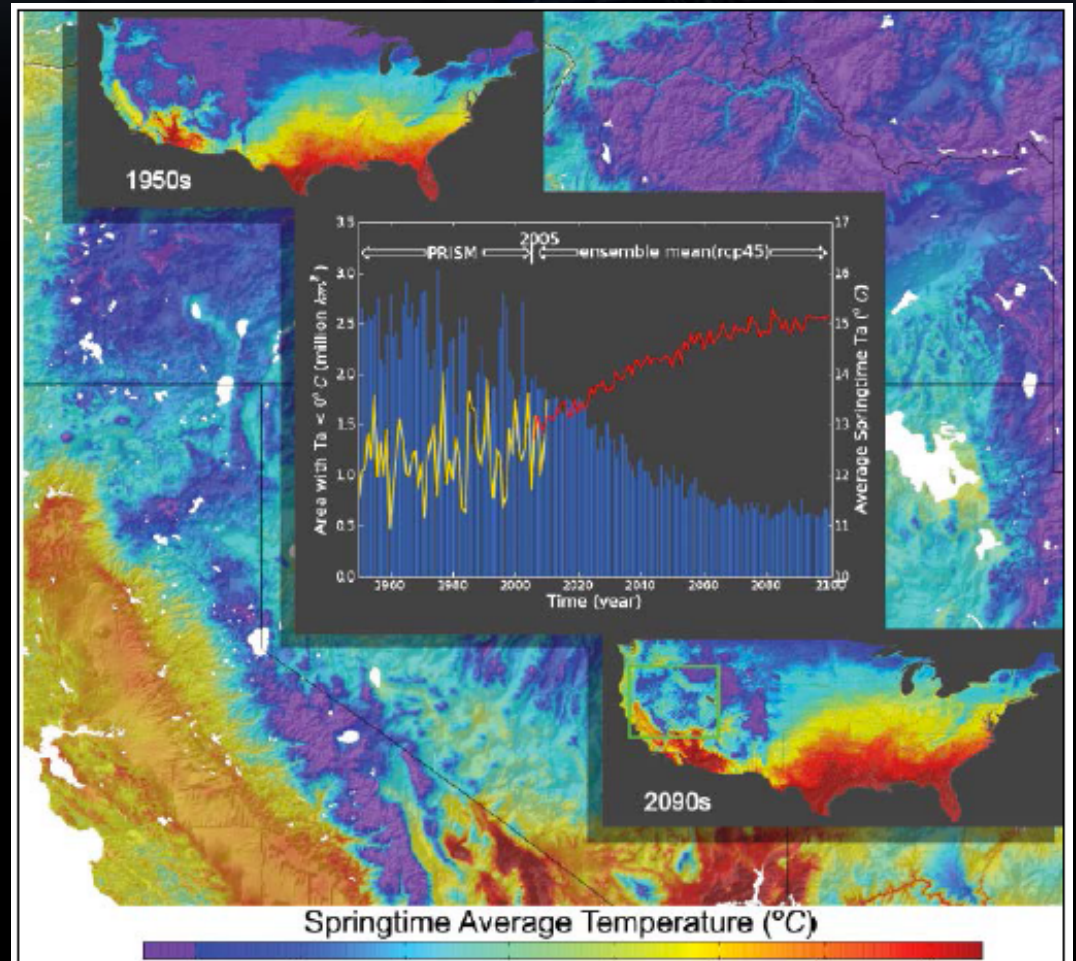


# High-Resolution Climate Projections



## National Climate Assessment

- Statistical downscaling of coarse data from CMIP5 (for IPCC) for conterminous U.S. to obtain high-resolution predictions at local scale
- ~800m grid resolution
- Spring (March–May), 1950–2099
- Mean temperature projected to increase from 12°C to 15°C assuming greenhouse gas emissions stabilize in 2050
- Area at or below 0°C isotherm decreases from 2.5M sq. km to 0.6M sq. km





- Details in paper by B. Thrasher et al.: *“Downscaled climate projections suitable for resource management,”* Eos, Vol. 94, No. 37, Sept. 2013, pp. 321-323

# Quantum Computing



- Quantum mechanics deals with physical phenomena at very small scales ( $\sim 100$  nm) or at very low temperatures (few K) where actions are quantized
- The outcome of a quantum experiment is probabilistically associated both with what was done before the measurement and how the measurement was conducted
- Qubits (quantum bits) can exist in a superposition of states, allowing  $n$  qubits to represent  $2^n$  states simultaneously
- At the end of a computation, on measurement, the system collapses into a classical state returning only one bit string as a possible solution







# Quantum Computing Implementations


**Nanoelectronics,  
NMR, Diamond  
Chips, ...**

## Trapped Ions and Trapped Neutral Atoms




2-qubit ion trap with  
microwave control  
(top); 300-qubit ion  
trap in optical lattice  
(bottom)  
(trapping and  
manipulation of ions  
and atoms)

## Photonic Quantum Chips



4-qubit photonic chip  
with optical  
waveguides  
integrated in solid  
state  
(position or  
polarization of  
photons used a  
qubit)



## Superconducting Qubits





4-qubit  
universal  
quantum  
computing

## D-WAVE "VESUVIUS"

512-qubit – not universal



RWTH Aachen 2-qubit gate



Quantum dots (top);  
spin states of  
molecules in liquid  
(middle); nitrogen  
vacancies in  
diamond (bottom)






# Quantum Annealing





A physical technique to solve combinatorial optimization problems in QUBO (Quadratic Unconstrained Binary Optimization)

$$E(z_1, z_2, \dots, z_n) = \underbrace{\left(1 - \frac{t}{T}\right) H_0(\{z\})}_{A(t)} + \underbrace{\frac{t}{T} H_P(\{z\})}_{B(t)}$$

- $N$ -bit string of unknown variables  $\{z\}$
- $H_0$ : Hamiltonian with known ground state
- $H_P$ : Hamiltonian whose ground state represents the solution to the problem
- $A(t)$  is slowly (adiabatically) lowered to zero while maintaining minimum energy of the system at all times
- Solution is the configuration  $\{z\}$  that produces the minimum  $E$  with a non-zero probability




# NASA and Quantum Computing



Data Analysis and  
Data Fusion

Anomaly Detection  
and Decision Making




Air Traffic  
Management


V&V and  
optimal  
sensor  
placement



Mission Planning and Scheduling, and Coordination



Topologically  
aware Parallel  
Computing






# D-Wave Two System



- Collaboration among NASA, Google, and USRA led to installation of system at NAS Division
- 512-qubit Vesuvius processor (to be continuously upgraded over the next 4+ years)
- 10 kg of metal in vacuum at 15 mK
- Magnetic shielding to 1 nanoTesla
- Protected from transient vibrations
- Single run takes 20  $\mu$ secs
- Uses 12 kW of electrical power

Focused on solving discrete optimization problems using quantum annealing





# Programming the D-Wave Two



## 1 Map the target combinatorial optimization problem into QUBO

No general algorithms, smart mathematical tricks (penalty functions, locality reduction..)


$$Q_{ijk} z_i z_j z_k$$

$$Q_{ijk} y_{ij} z_k + \beta_{ijk} (3y_{ij} - 2z_i y_{ij} - 2z_j y_{ij} + z_i z_j)$$

$$\sum_{ij} Q_{ij} z_i z_j$$

## 2 Embed the QUBO coupling matrix in the hardware graph of interacting qubits


The D-Wave hardware qubit connectivity is a “Chimera Graph”, so embedding methods mostly based on heuristics



Note: D-Wave provides a heuristic blackbox compiler that bypasses embedding

## 3 Run the problem many times and collect statistics

Use symmetries, permutations, and error correction to eliminate the systemic hardware errors and check the solutions




*Mapping not needed for random spin-glass models*








*Embedding not needed for native Chimera problems*

*Performance can be improved dramatically with smart pre-/post-processing*

# Advanced Computing Mission



*Enable the science & engineering required to meet NASA's missions and goals*




*Effective, stable, production-level HPC environment*



*Advanced technologies to meet future goals*





# Thank You!

## Questions?

[rupak.biswas@nasa.gov](mailto:rupak.biswas@nasa.gov)